

**DATA WAREHOUSE APPLICATIONS FOR
NETWORKS OF SELF-SERVICE MACHINES**

Background of the Invention

5 1. **Field of the Invention**

The present invention relates to systems and methods of data warehousing applications, and in particular to data warehouse applications for networks of self-service machines (SSMs).

2. **Description of the Related Art**

10 Database management systems are used to collect, store, disseminate, and analyze data. Large-scale, integrated database management systems provide an efficient, consistent, and secure data warehouse capability for storing, retrieving, and analyzing vast amounts of data. This ability to collect, analyze, and manage massive amounts of information has become a virtual necessity in business today.

15 The information stored by these data warehouses can come from a variety of sources. One important data warehouse application involves the collection and analysis of information collected in the course of transactions between businesses and consumers. For example, when an individual uses a debit or credit card to access an automated teller machine (ATM) or purchase an item at a retail store, the identity of the customer, the transaction, and other
20 related information are collected. Traditionally, this information is used only to determine if the transaction should be completed, and rarely for other purposes.

However, such data could also be used to identify demographic information concerning the customers and cross-sell other services and products. Unfortunately, the infrastructure of SSM networks does not support such uses of the information. Thus, there is
25 a need in the art for mechanisms that allow for the exploitation of this information.

Summary of the Invention

To overcome the limitations in the prior art described above, and to overcome other limitations that will become apparent upon reading and understanding the present

specification, the present invention discloses a network of interconnected self-service machines (SSMs), wherein each of the SSMs executes a relational database management system (RDBMS) that maintains a relational database stored on the SSM. The relational databases are partitions of a global relational database and each partition stores information for only those customers that frequent the SSM that executes the RDBMS. The relational databases are used to more effectively serve the customer at the SSM and to market products and services to the customer at the SSM.

Brief Description of the Drawings

Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

FIG. 1 illustrates an exemplary hardware environment that could be used to implement the preferred embodiment of the present invention;

FIG. 2 further illustrates the exemplary hardware environment of FIG. 1;

FIG. 3 illustrates the logic performed by the data warehouse system according to the preferred embodiment of the present invention; and

FIG. 4 illustrates the logic performed by the self-service machines according to the preferred embodiment of the present invention.

Detailed Description of the Preferred Embodiment

In the following description of the preferred embodiment, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration a specific embodiment in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

System Description

FIG. 1 illustrates an exemplary hardware environment that could be used to implement the preferred embodiment of the present invention. In the exemplary environment, a network 100 interconnects a plurality of self-service machines (SSMs) 102,

wherein the SSMs 102 may comprise automated teller machines (ATMs), computer-based retail kiosks, computer-based product vending machines, and the like. The network 100 also connects the SSMs 102 with one or more transaction processing systems 104 to process transactions, such as informational, retail, commercial, banking, financial or other transactions. One or more data warehouse systems 106 are also connected to the network 100 for communication with both the SSMs 102 or transaction processing systems 104, wherein the data warehouse system 106 collects, stores, analyzes, and disseminates information collected in the course of financial or commercial transactions between the SSMs 102 and the transaction processing system 104.

FIG. 2 further illustrates the exemplary hardware environment of FIG. 1. In this illustration, each the SSMs 102 are shown to comprise one or more processors 200 coupled to one or more fixed and/or removable data storage units (DSUs) 202, such as disk drives, that store data and programming. Each of the SSMs 102 operate under the control of an operating system (OS) 204 and executes a relational database management system (RDBMS) 206 that maintains a relational database 208 stored on the DSUs 202.

SSMs 102 may be operated in large networks. For example, large financial institutions may have between 2,000 and 8,000 ATMs. These financial institutions also are beginning to build centralized data warehouse systems 106 in order to collect information about their customers and to use that information to increase propensity, profitability and customer satisfaction.

However, using a centralized data warehouse system 106 to provide information-rich customer interaction at SSMs 102 is currently prohibitive, since it would likely lengthen the amount of time involved in the SSM 102 transaction. For example, when using an SSM 102, a customer inserts their card into the SSM 102, their pin is validated locally within the SSM 102, a limited menu selection allows the customer to create the transaction request, and then the transaction request is transmitted to the transaction processing system 104 for authorization.

This idle time is ideal for cross-selling opportunities. If this idle time is only slightly increased to retrieve and display information for the customer, no harm is done.

Unfortunately, if the idle time is substantially increased, or perceived to be increased, substantial harm may be done to the relationship with the customer.

Generally, the customer just wants to complete their transactions and leave. They are often unwilling to wait to absorb a cross-selling message after the transaction completes.

5 Moreover, if the customer is forced to remain at the SSM 102 to receive the cross-selling message, the overall transaction time rises, the SSM 102 queue gets longer, and customer service is perceived to suffer.

The majority of SSM 102 customers are creatures of habit and interact with a small number of SSMs 102 at a small number of predictable locations for the majority of their
10 transactions (this is especially true with ATMs, for example). As indicated in FIG. 2, many SSMs 102 comprise sophisticated computer systems; indeed, most will be based on sophisticated operating systems 204, such as WINDOWS NT™, within the next few years. As a result, more sophisticated customer-specific programming can be located at the SSMs 102 themselves rather than rely on access to the centralized data warehouse system 106 with
15 its incumbent delays.

The preferred embodiment of the present invention essentially transforms the SSMs 102 into sophisticated platforms for RDBMS 206 applications that structure information, transactions, services, product offerings and the like based on the customer information maintained by the RDBMS 206. Moreover, the network 100 of interconnected SSMs 102
20 can be considered a massively parallel processor (MPP) computer system that operates parallelized RDBMS's 206. The basic intent of the present invention is to view the individual relational databases 208 as subsets or partitions of a global relational database, and each of the individual RDBMS's 206 as comprising part of a global RDBMS.

In the preferred embodiment of the present invention, the OS 204 comprises the
25 WINDOWS NT™ product offered by Microsoft Corporation, although other OS's could be used as well. Moreover, in the preferred embodiment of the present invention, the RDBMS 206 software comprises the Teradata® product offered by NCR Corporation, which is designed specifically for operation on MPP computer systems, although other RDBMS's could be used as well.

The global relational database is divided among the SSMs 102 by storing customer-specific information at specific SSMs 102. Thus, each DSU 202 stores a subset of the global relational database and operations can be performed either locally or globally. In the preferred embodiment, the subset includes information on those customers that frequent the particular SSM 102.

Because the information is stored locally on the SSM 102, the SSM 102 is capable of using information about the customer, at the point of interaction, to more effectively serve the customer, to provide personalized service, as well as to market products and services to the customer. Moreover, because the RDBMS 206 is a standard software product, it is available to other SSMs 102 or other systems on the network 100.

It is one goal of the present invention to parallelize functions across the SSMs 102. For example, queries or other operations can be divided into smaller sub-queries, each of which can be assigned to an SSM 102 based on the specific information stored on the SSMs 102. Alternatively, specific operations on particular information can be directed to an SSM 102 storing the particular information.

In essence, the present invention decentralizes the centralized data warehouse system 106 to the network 100 of SSMs 102. Note that the network 100 of SSMs 102 may or may not be considered the primary data warehouse system; it may be just a duplicate or back-up of the data warehouse system 106. Moreover, the SSMs 102 and the centralized data warehouse system 106 could be synchronized at night or during other low usage periods. Further, the SSMs 102 may synchronize or move data among each other as required.

The SSMs 102 may also capture detailed data about the customer's interaction (screens viewed, information presented, dwell time, and transaction information, if any) for use both locally and globally. This detailed data could be uploaded to populate the centralized data warehouse system 106 or copied to another SSM 102.

Based on customer usage patterns as determined by the data warehouse system 106, data may be stored on multiple SSMs 102 or moved among SSMs 102 as required. Generally, customer usage patterns are fairly constant, and thus only a minimal amount of data is moved at any time.

Such a system ties together the SSMs 102 and the centralized data warehouse system 106 in interesting and powerful ways. It demonstrates the leverage and utilization that can be applied to a generally underutilized SSM 102 infrastructure. Moreover, it provides a foundation for more sophisticated applications.

Logic of the Data Warehouse System

FIG. 3 illustrates the logic performed by the data warehouse system 106 according to the preferred embodiment of the present invention.

Block 300 represents the data warehouse system 106 collecting information from the SSMs 102 and/or the transaction processing systems 104 in the course of financial or commercial transactions between the SSMs 102 and the transaction processing systems 104.

Block 302 represents the data warehouse system 106 analyzing the collected information to identify the locations at which specific customers interact with the SSMs 102.

Block 304 represents the data warehouse system 106 downloading customer-specific information to the identified SSMs 102. This step may also include downloading customer-specific information retrieved from the transaction processing system 104.

Logic of the Automated Teller Machine

FIG. 4 illustrates the logic performed by the SSMs 102 according to the preferred embodiment of the present invention.

Block 400 represents the SSM 102 retrieving customer-specific information from the relational database 208 stored on the SSMs 102. This information can be retrieved at any time after the customer is identified by the SSM 102. Alternatively, this step may require the retrieval of customer-specific information from the transaction processing system 104, the data warehouse system 106, or one or more of the other SSMs 102, especially if the SSM 102 is not one frequented by the customer. The customer-specific information is used to more effectively serve the customer, to provide personalized service, as well as market products and services to the customer.

Block 402 represents the SSM 102 generating a customer-specific display on the SSM 102. This step may be performed before or after any customer interaction (after the customer has been identified, of course). This step may also be performed during any idle

time while the customer is waiting for completion of a transaction or information request. The goal is not to increase the amount of time it takes for the transaction to complete.

Block 404 represents the SSM 102 processing a customer transaction or request. For example, this step may include accepting a request for information, services, or products from the customer into the SSM 102, validating the transaction (either locally or remotely via the transaction processing system 104), presenting additional menu selections, and then communicating with the transaction processing system 104 to complete the transaction.

Conclusion

This concludes the description of the preferred embodiment of the invention. The following describes an alternative embodiment for accomplishing the same invention. In one alternative embodiment, any type of computer or network could be used to implement the present invention. In addition, any database management system could benefit from the present invention.

In summary, the present invention discloses a network of interconnected automated teller machines (SSMs), wherein each of the SSMs executes a relational database management system (RDBMS) that maintains a relational database stored on the SSM. The relational databases are partitions of a global relational database and each partition stores information for only those customers that frequent the SSM that executes the RDBMS. The relational databases are used to more effectively serve the customer at the SSM and to market products and services to the customer at the SSM.

The foregoing description of the preferred embodiment of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto.